

# FastTrans: A Scalable and Parallel Discrete Event Microsimulator for Transportation Networks

Sunil Thulasidasan, CCS-3;  
Shiva Kasiviswanathan, IBM;  
Stephan Eidenbenz, CCS-3; Phil Romero, HPC-1

FastTrans is a scalable, parallel microsimulator for transportation networks that can simulate and route tens of millions of vehicles on real-world road networks in a fraction of real time. FastTrans uses parallel discrete-event simulation techniques and distributed-memory algorithms to scale simulations to over one thousand compute nodes.

Traditionally, traffic microsimulations of transportation networks have employed a time-stepped cellular-automata approach. A prominent example of this is TRANSIMS, developed at LANL, where vehicular dynamics are modeled at a high level of spatial granularity. This allows one to capture phenomena such as lane changing and vehicular emissions, but comes at a high computational cost. A discrete-event queue-based model for transportation networks was first proposed in [1], but this was limited to a sequential, single-processor environment.

*Fig.1. A\* explores only a fraction of nodes (blue) compared to Dijkstra (green) to reach the destination.*



The FastTrans approach is to combine the discrete-event queue model with scalable parallelization. This allows us to simulate large-scale, real-world networks and realistic traffic scenarios involving tens of millions of vehicles in a fraction of real time. Also, since FastTrans simulates the behavior of each vehicle or traveler at the individual entity level, it retains some of the advantages of microsimulations. In addition, the congestion feature implemented in FastTrans, which updates the state of the routing graph on all simulation processes, allows one to observe the macroscopic behavior of the network.

Vehicular trips are generated using agent-based simulations that provide

realistic, daily activity schedules for a synthetic population of millions of intelligent agents. FastTrans can execute simulations of large cities up to 20 times faster than real time as a result of utilizing a queue-based approach to road network modeling, which has been shown to be significantly faster than traditional approaches based on cellular automata models—at the same time it captures road link and intersection dynamics with high fidelity.

The routing algorithm, which is the most computationally intensive part of FastTrans, is a heuristic search variant of the classic Dijkstra shortest-path algorithm (A\*). FastTrans uses a highly optimized version of A\* that uses the structural properties of the road network and performs a goal-directed search to find the best path toward the destination. Experimental results have shown up to 30-fold improvements [2,3] in routing performance compared to naive implementations of shortest-path algorithms. FastTrans is also able to achieve near-perfect load-balancing on distributed clusters, through an optimal assignment of simulation entities to processors using a technique called explicit spatial scattering [4] (see Figs. 1-3).

With optimized routing and partitioning, FastTrans is able to simulate a full 24-hour work day in New York—involving over one million road links and approximately 25 million vehicular trips—in less than one hour of wall-clock time on a 512-node cluster. The quick turn-around capability of FastTrans has been employed in numerous exercises for the Department of Homeland Security, in which infrastructure disruptions in areas like New York City, Southern California, and the Twin Cities were

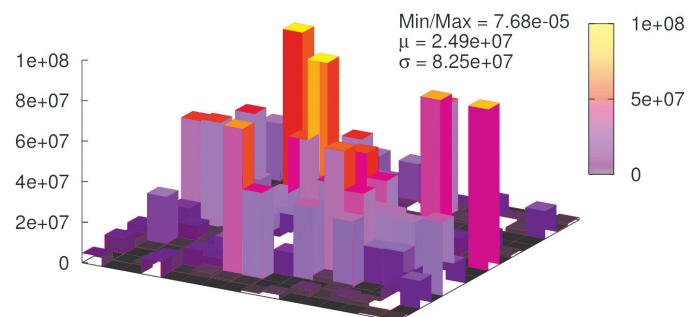


Fig. 2. Traditional load balancing schemes result in uneven load distribution.

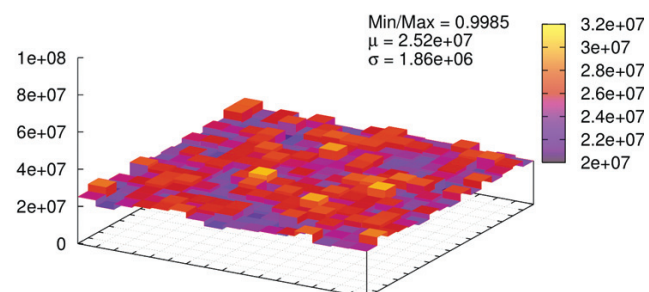


Fig. 3. Scatter partitioning results in evenly distributed load.

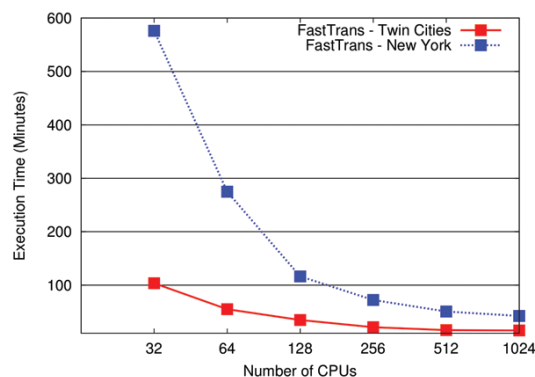


Fig. 4. Performance scaling with increasing cluster sizes.

simulated and analyzed. Recent capability additions to FastTrans include the ability to simulate multi-modal transportation networks; future capability augmentations include hierarchical routing for whole-nation studies, and multi-scale simulations on hybrid architectures (see Fig. 4).

- [1] Charypar, initials, et al., "An Event-driven Queue-based Microsimulation of Traffic Flow," in *Transport Systems Planning and Transport Telematics* (2006). (is this a book? Need Eds, publisher)
- [2] Thulasidasan, S., and S. Eidenbenz, "Accelerating Traffic Microsimulations: A Parallel Discrete-Event Queue-based Approach for Speed and Scale," in *Proc ACM/SIGSIM Winter Sim Conf* (2009).
- [3] Thulasidasan, S., et al., "Designing Systems for Large-Scale, Discrete-Event Simulations: Experiences with the FastTrans Parallel Microsimulator," in *Proc Intl Conf High Perform Comput (HiPC)* (2009).
- [4] Thulasidasan, S., et al., "Explicit Spatial Scattering for Load Balancing in Distributed Discrete Event Simulations with Conservative Time Synchronization," in *Proc ACM/IEEE/SCS Workshop Principles Adv Distrib Sim* (2010).

## Funding Acknowledgments

National Infrastructure Simulation and Analysis Center (NISAC) of Department of Homeland Security